

Denial and Deception in Cyber Defense

Kristin E. Heckman, Frank J. Stech, Ben S. Schmoker, and Roshan K. Thomas, The MITRE Corporation

As attack techniques evolve, cybersystems must also evolve to provide the best continuous defense. Leveraging classical denial and deception techniques to understand the specifics of adversary attacks enables an organization to build an active, threat-based cyber defense.

It is now widely recognized that traditional approaches to cyber defense have been inadequate. Boundary controllers and filters such as firewalls and guards, virus scanners, and intrusion detection and prevention technologies have all been deployed over the last decade. Nevertheless, sophisticated adversaries using zero-day exploits are still able to get in and, in many cases, establish a persistent presence. We must assume, then, that an adversary will breach border controls and establish footholds within the defender's network, so we need to study and engage the adversary on the defender's turf in order to influence any future moves. A key component in this new paradigm is cyber denial and deception (cyber D&D).

The goal of D&D is to influence another to behave in a way that gives the deceiver an advantage, creating a causal relationship between psychological state

and physical behavior. *Denial* actively prevents the target from perceiving information and stimuli; *deception* provides misleading information and stimuli to actively create and reinforce the target's perceptions, cognitions, and beliefs. Both methods generate a mistaken certainty in the target's mind about what is and is not real, making the target erroneously confident and ready to act.

As adversaries' attack techniques evolve, defenders' cybersystems must also evolve to provide the best continuous defense. Engineering cybersystems to better detect and counter adversarial D&D tactics and to actively apply D&D against advanced persistent threats will force adversaries to move more slowly, expend more resources, and take greater risks. In doing so, defenders may possibly avoid, or at least better fight through, cyber degradation.

D&D TECHNIQUES

Table 1 shows a two-dimensional framework to apply D&D techniques.¹ The first dimension relates to information (fact or fiction); the second relates to actions or behaviors (revealing or concealing). The deceiver uses denial to prevent the detection of the *essential elements of friendly information* (EEFI) by hiding what's real, and employs deception to induce misperception by using the *essential elements of deception information* (EEDI) to show what's false. The deceiver also has to hide the false information—that is, the *nondisclosable deception information* (NDDI)—to protect the D&D plan, and additionally show the real information—the *nonessential elements of friendly information* (NEFI)—to enhance the D&D cover story. Deception is a very dynamic process, and deception planners will benefit from the interplay of techniques from more than one quadrant in a deception operation.

Cyber D&D maps this framework to the cyber domain. Table 2 uses the D&D methods matrix to show high-level cyber D&D techniques (combinations of two or more tactics), organized according to whether they are fact or fiction, and whether they are revealed via deception methods or concealed via denial methods. This “packaging” of tactics can be an indicator of a sophisticated deception capability maturity (CM) for counterdeception (CD) purposes.

THE DECEPTION CHAIN

The *deception chain* is a high-level meta model for cyber D&D operations management from a life-cycle perspective. Analogous to Lockheed Martin's “cyber kill chain” model,² the deception chain is adapted from Barton Whaley's 10-step process for planning,

TABLE 1. D&D methods matrix.

Deception objects	Deception: Mislead (M)-type methods	Denial: Ambiguity (A)-type methods
	Revealing	Concealing
Facts	Reveal facts: Nonessential elements of friendly information Reveal true information to the target Reveal true physical entities, events, or processes to the target	Conceal facts (dissimulation): Essential elements of friendly information Conceal true information from the target Conceal true physical entities, events, or processes from the target
Fiction	Reveal fiction (simulation): Essential elements of deception information Reveal false information to the target Reveal false physical entities, events, or processes to the target	Conceal fiction: Nondisclosable deception information Conceal false information from the target Conceal false physical entities, events, or processes from the target

TABLE 2. D&D methods matrix with cyber D&D techniques.

Deception objects	Deception: M-type methods	Denial: A-type methods
	Revealing	Concealing
Facts	Reveal facts: Nonessential elements of friendly information Publish true network information Allow disclosure of real files Reveal technical deception capabilities Reveal misleading, compromising details Selectively remediate intrusion	Conceal facts (dissimulation): Essential elements of friendly information Deny access to system resources Hide software using stealth methods Reroute network traffic Silently intercept network traffic
Fiction	Reveal fiction (simulation): Essential elements of deception information Misrepresent intent of software Modify network traffic Expose fictional systems Allow disclosure of fictional information	Conceal fiction: Nondisclosable deception information Hide simulated information on honeypots Keep deceptive security operations a secret Allow partial enumeration of fictional files

preparing, and executing deception operations.³ The deception chain facilitates the integration of three systems—cyber D&D, cyber intelligence, and security operations—into the enterprise's larger active defense system to plan, prepare, and execute deception operations.

Deception operations are conducted by a triad of equal partners

working those three systems inter-actively: cyber D&D planners, cyber-intelligence analysts, and cyberse-curity operators. Just as computer network defense (CND) is not any one tool but a system that deploys new technologies and procedures as they become available, cyber D&D must be thought of as an active defensive oper-ational campaign, employing evolving

tools, tactics, techniques, and procedures (TTTPs). We believe the deception chain is a flexible framework for embedding advanced TTTPs in operational campaigns while focusing on an organization's mission objectives.

This triad of cyber D&D planners, cyber-intelligence analysts, and cybersecurity operators is essential for a threat-based active defense. There are eight phases in the deception chain.

Purpose

This initial phase helps enterprise managers define the strategic, operational, or tactical goal—in other words, the purpose of the deception—and the criteria that would indicate the deception's success.

Collect intelligence

In the next phase, D&D planners define how the adversary is expected to react to the deception operation. This is done in part through the planners' partnership with cyber intelligence to determine what the adversary will observe, how the adversary might interpret those observations, how the adversary might react (or not) to those observations, and how to monitor the adversary's behavior. This intelligence will help planners during the last two phases (monitor and reinforce) to determine whether the deception is succeeding. One internal source of intelligence is *intrusion campaign analysis*.⁴ Broadly speaking, an intrusion campaign is a framework that combines all the related information about a particular intrusion into a set of activities.

Threat-sharing partnerships are another source of cyber intelligence and might involve government, private industry, or nonprofit organizations. For example, the Structured Threat

Information eXpression (STIX; <http://stix.mitre.org>) and Trusted Automated eXchange of Indicator Information (TAXII; <http://taxii.mitre.org>) systems, sponsored by the Office of Cybersecurity and Communications at the US Department of Homeland Security, provide a structured format for defenders to share threat indicators in a manner that reflects the trust relationships inherent in such transfers. STIX is a community-driven language used to represent structured cyber-threat information, and TAXII enables the sharing of information across organization and product boundaries to detect and mitigate cyber threats.

A threat seen by one partner today might be the threat facing another partner in the near future. All of these sources of cyber intelligence can aid D&D planners in assessing an adversary's cyber D&D CM, which in turn supports the development of an appropriately customized cyber D&D operation.

Design cover story

The *cover story* is what the defender wants the adversary to perceive and believe. The D&D planner will consider the critical components of the operation, assess the adversary's observation and analysis capabilities, and develop a convincing story that "explains" the operation's components observable to the adversary but misleads the adversary as to the meaning and significance of those observations.

The D&D planner will decide what information must be hidden (the EEFI and the NDDI) and what information must be created and revealed (the EEDI and the NEFI). The D&D methods matrix in Table 1 aids planners by capturing the true and false information that must be revealed or concealed to make the deception operation

effective. The planners and cybersecurity operators must decide what information "belongs" in the four cells of the matrix and get buy-in from enterprise managers for the deception goals and cover story.

Plan

In this phase, D&D planners analyze the characteristics of the real events and activities that must be hidden to support the deception cover story, identify the corresponding signatures that would be observed by the adversary, and plan to use denial tactics (such as masking, repackaging, dazzling, or red flagging³) to hide the signatures from the adversary. Planners also analyze the characteristics of the notional events and activities that must be portrayed and observed to support the cover story, identify corresponding signatures the adversary would observe, and plan to use deception tactics (such as mimic, invent, decoy, or double play³) to mislead the adversary.

In short, D&D planners turn the matrix cell information into operational activities that reveal or conceal the key information conveying the cover story. These steps must be coordinated with security operations so that they are as realistic and natural as possible, and the deception should allow the adversary to observe real operational events that support the cover story.

Prepare

In this phase, D&D planners design the desired effect of the deception operation and explore the available means and resources to create the effect on the adversary. This entails coordination with security operations on the timing for developing the notional and real equipment, staffing, training,

MALICIOUS ACTORS FOLLOW A COMMON MODEL OF BEHAVIOR TO COMPROMISE VALUABLE INFORMATION IN A TARGET NETWORK.

and other preparations to support the deception cover story.

Execute

If the deception and real operational preparations can be synchronized and supported, the D&D planners and security operations must coordinate and control all relevant preparations so they can consistently, credibly, and effectively execute the deception cover story.

Monitor

D&D planners work with cyber intelligence and security operations to monitor and control the deception and real operations. This entails monitoring both friendly and adversary operational preparations, carefully watching the observation channels and sources selected to convey the deception to the adversary, and monitoring the adversary's reaction to the "performance," that is, the cover story execution. These targeted channels must remain open to the adversary, convey the planned deception, and be observed by the adversary.

Reinforce

If cyber intelligence on the adversary indicates that the deception operation does not seem to be "selling" the cover story to the adversary, the D&D planners may need to reinforce the cover story through additional deceptions, or to convey the deception operation to the adversary through other channels or sources. The planners may have to revisit the first phase of the deception chain, execute a backup deception, or plan another operation.

THE DECEPTION CHAIN AND THE CYBER KILL CHAIN

Malicious actors follow a common model of behavior to compromise

valuable information in a target network. Attackers generally employ a cyberattack strategy, divided into the six phases described below, called the cyber kill chain or kill chain.² Like the cyber kill chain, the deception chain is not always linear. Progression through the phases can be recursive or disjoint. One run of the kill chain models a single intrusion, but a campaign spanning multiple engagements builds on previous results and omits phases as necessary. Similarly, D&D planners and cybersecurity operators will selectively run through the deception chain to achieve their goals.

The deception chain is also applicable at each phase of the cyber kill chain, as illustrated by the following hypothetical deception operation goals associated with each kill chain phase. Later, we describe two case studies that will further demonstrate the interplay of the kill chain and the deception chain to defend against intrusions. Phase names in the case studies are italicized to emphasize the interconnectedness of the two chains.

- › *Recon*: If defenders are aware of adversarial reconnaissance efforts, provide the adversary with a set of personae and a Web footprint for defensive targeting efforts in the delivery phase. Note that deception operations can be used to influence adversary actions in future kill chain phases.
- › *Weaponize*: Making the adversary (wrongly) feel certain about an organization's vulnerabilities, defense posture, or capabilities could enable the organization to recognize or defend against the adversary's weaponized payload. If the recon phase was successful, the adversary

will attempt to deliver the weaponized payload to one or more of the false personae.

- › *Exploit*: Recognizing exploitation attempts, defenders may redirect the adversary to a *honeypot* environment, which appears to be part of a network that contains valuable information but is actually isolated and monitored by defenders. The goal is to conceal all honeypot "tells" or indicators to delay the adversary.
- › *Control*: When the adversary has "hands on keyboard" access, provide the adversary with a high-interaction honeypot with a rich variety of information, designed with the D&D planners, to help identify the adversary's motives, intentions, and capability maturity.
- › *Execute*: Slow the adversary down to collect cyber intelligence.
- › *Maintain*: Keep up the appearance of realism in a high-interaction honeypot by adding or retiring false personae, as well as maintaining existing personae and their "pocket litter," such as files, email, password change history, login history, and browser history.

These examples also show that there may be a need for more than one deception operation during a single intrusion.

CASE STUDY: STOPPING A "SMASH AND GRAB" INTRUDER

In 2013 the Syrian Electronic Army (SEA) attacked a number of news agencies, compromising user accounts and

defacing public websites via stolen access. This group prioritized speed over subtlety, sending hundreds of emails eliciting user account credentials.⁵ The attacks were eventually mitigated with an organization-wide password reset along with second-factor authentication for sensitive information. Victims in this case could have leveraged D&D to both prevent compromise and allow faster response to the attack.

To foil the *recon* phase of this intrusion, nonexistent employees with plausible email addresses could be placed on public-facing websites. These mailboxes serve to solicit unwanted messages and proactively notify network defenders of attempts to target publicly visible employees. Defenders would *prepare* these false targets while *reinforcing* their effectiveness, *executing* on the attacker's intrusion attempt, and *monitoring* the results.

The attackers gained remote access to internal user accounts, catapulting them past the *exploit* phase. Targeting specific information rather than trying to maintain long-term access during the *control* phase, the SEA quickly sent another round of internal emails soliciting access to the content production network. Several hours later, news headlines were modified with propaganda from the attackers.

Implementing an internally hosted honeypot to emulate the news agency's content production system may have alerted security teams to the SEA's attack. Because the only *purpose* of these systems is to collect suspicious requests, it should be assumed that any connections to these systems are malicious.

Another way to detect attackers who are consolidating and expanding their access involves seeding internal

systems with tripwire user accounts. By creating plausible "administrator" or "maintenance" accounts, defenders can know a system is compromised as soon as the login attempt is denied.

The SEA's goal was to modify news content during the *execute* phase. At this point in the attack, defenders have to assume that an attacker is as threatening as a malicious insider. Wary defenders can attract attempts to access sensitive information by creating tantalizing internal documents that are visible but not accessible. In this case, the news agency could have prepared headlines that an adversary would find interesting, and log any attempts to view or modify that internal data.

CASE STUDY: MITIGATING A "LOW AND SLOW" INTRUDER

Another intrusion group, active since at least 2006 and identified as APT1 by investigators in 2013, used tactical D&D to prevent detection and achieve their mission.⁶ Over the course of months and years, their victims suffered intellectual property theft and system compromises.

During the *exploit* phase, APT1 attempted to social-engineer individuals with malicious emails from spoofed senders related to the target's area of interest. APT1 crafted and attached documents to emails to exploit targets, which avoided email scanning tools by appearing legitimate. Once they were able to *execute* on an infected system, their malware mimicked legitimate system services to evade detection.

The attackers leveraged a blend of publicly available malware and custom tools to *control* infected systems using a network of proxies and compromised servers. Several tools implemented

covert channels that mimicked client traffic such as that from chat and Web services, further decreasing the chance of detection. To *maintain* access to target networks, administrative credentials were used to establish an entrenched foothold and obtain sensitive information.

Given that APT1's objective was to collect valuable data from a large number of targets over an extended period of time, future mitigation efforts would need to apply to a broad range of companies and government agencies.⁶ One mitigation strategy might be to *prepare* for intrusions by building false assets that appear realistic and re-engineering real assets to look like decoys.

After APT1 had compromised a target, investigators could *monitor* activity related to those assets and share lessons learned with partner organizations. Threat intelligence on the attacker's behavior and tools could then be used to *reinforce* the efficacy of future deception operations. Defenders would leverage this intelligence and be better able to *execute* a plan whose *purpose* is to influence and manipulate where and how the adversary can operate.

TOWARD A CYBER D&D MATURITY MODEL

Like any other capability introduced into an organization, cyber D&D must be carefully coordinated and managed to achieve the desired results. The most critical facets are the maturity level and the overall management model.

The cyber D&D maturity level manifests in the people, services, processes, and technologies specifically enabled to conduct cyber D&D operations; key indicators and metrics involve the degree to which cyber D&D capabilities are systemized and optimized. A cyber

D&D capability maturity model provides a coherent blueprint that organizations can use to assess, measure, and increase the maturity of their current cyber D&D operations and develop specific cyber D&D innovations. It can also help organizations characterize not only their own cyber D&D CM, but also that of their adversaries by providing observable and often measurable indicators of specific cyber D&D capabilities. For example, organizations with a level 1 initial CM have ad hoc and chaotic processes that can be readily anticipated and countered and so fail to mislead the adversary or to attain the deception objectives; they lack the ability to correctly characterize an adversary's response. In contrast, organizations with a level 4 or 5 CM have processes that are repeatable, customized to individual adversaries, and not obvious. They also incorporate interdomain data and anticipate the adversary's response to deceptions.

The cyber D&D management model represents the overall approach to managing cyber D&D capabilities and operations from the perspectives of capability and operations and services. The operational processes used to conduct cyber D&D operations constitute another salient aspect of life-cycle management. In particular, cyber D&D must function in concert with the organization's overall defensive operations and must support cyber defense. The preparations undertaken to launch and manage D&D capability must encompass coordination among people, services, processes, and technology development and deployment.

The last two facets of life-cycle management include observables and indicators as well as related metrics. These observables and indicators vary with the life-cycle of D&D operations,

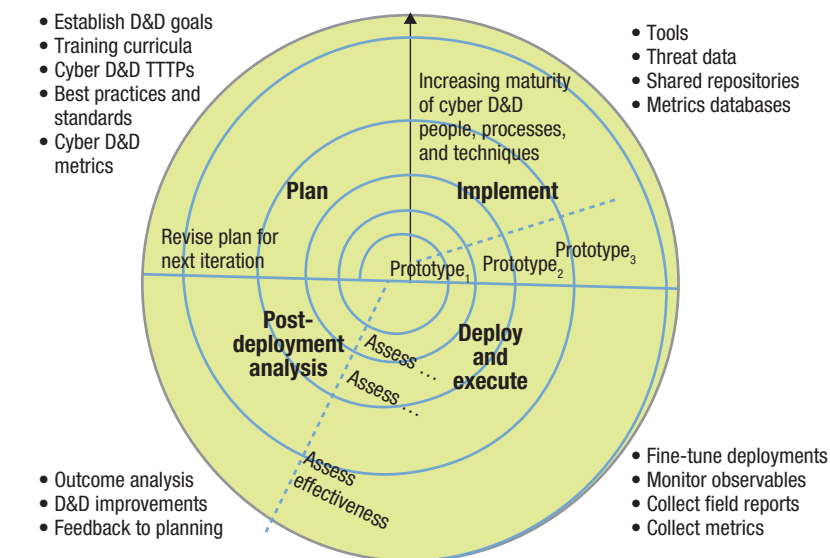


FIGURE 1. Overview of a spiral cyber denial and deception life-cycle management process. The four stages (plan, implement, deploy and execute, and post-deployment analysis) are conducted iteratively with an assessment of the effectiveness of each stage and a revision plan after each iteration. This life-cycle support model increases the maturity of an organization's cyber D&D capability.

but should give the cyber D&D team insights into the progress and effectiveness of its activities. Organizations must establish a metrics framework to quantitatively and qualitatively measure and track the various observables and indicators, and to run analytics on them to enable higher-level inferences about adversary TTTPs.

Organizations may benefit from a spiral D&D life-cycle management process that iteratively drives toward increasing overall cyber D&D capability effectiveness through higher maturity and continuous process improvements (see Figure 1). This type of model helps an organization assess risks and effectiveness with each iteration of the spiral while promoting agile and rapid prototyping as well as tuning of D&D techniques and services based on observed outcomes.

Any attempt to incorporate cyber D&D into active cyber defense must start with establishing clear and achievable program goals. At a minimum, the planning phase should include establishing D&D program goals and developing training curricula, cyber D&D TTTPs, cyber D&D best practices and standards, and cyber D&D metrics.

In the *implementation* phase, the organization will start to plan based on the goals and actions from the previous phase. The plan must address both the “what” (what artifacts are needed, such as tools and code bases, threat data, shared repositories, and metrics databases) and the “how.”

Next, the organization must *deploy* and *execute* cyber D&D TTTPs, services, and supporting processes in a target environment. The target environment could be a synthetic environment such as a honeynetwork or a honeypot, the real cyber infrastructure, or some combination. In accordance with the spiral methodology, this approach to life-cycle management calls upon the organization implementing the cyber D&D program to iterate through rapid prototypes of cyber D&D with each spiral. At each iteration, the organization must evaluate the risks and effectiveness of the current prototype.

Post-deployment analysis, the last phase in the spiral, has three essential elements: outcome analysis, process improvements, and feedback. Outcome analysis centers on the overall outcome of the current spiral, addressing questions such as:

- › How effective were the cyber D&D techniques developed and operationally deployed?
- › What were the successes and failures?
- › How well did the organization manage the total life-cycle costs within the spiral?

To answer these questions, the organization must analyze metrics data and field reports, using the results to formulate specific D&D improvements in processes, services, and technologies. This requires careful attention to managing change for all of the D&D elements. This phase also generates feedback to help with the next iteration of the spiral model. The organization must share appropriate metrics data with relevant parties—those refining D&D techniques or those planning the next iteration.

A system of deceptive interactions between intruders and defenders benefits from both technological development and operational coordination with cyber-intelligence analysts and security operators.

Is there a technology development roadmap that can guide academic and commercial developments? These questions are beyond the scope of this article, but we propose some research and technology challenges and present them within the framework of the deception chain.

To inform the purpose and collect intelligence stages, models for strategic D&D objectives can be built from both offensive and defensive perspectives. An organization that understands adversarial models can more quickly tailor defensive operations to their adversaries' capabilities and intent. These actions may be tactical and operational as well as strategic. Game-theory models could help analyze moves and countermoves to produce promising TTTPs for cyber D&D.

During the *cover story* stage, it is important to create believable deception material to attract the adversary's interest. Network and host-based deception material such as honeypots, crafted documents, and email are referred to as *honeytokens*. These are currently crafted manually by

To what extent do these moves signal the adversary's intentions? Which baits have worked well, or not? What is the adversary's sphere of interest?

Preparing and executing a deception operation can be made more scalable and efficient by leveraging existing tools and training materials. A standalone "honeypot in a box" product might be developed to adapt to an organization's network structure with a truncated setup time. Novel ways of training personnel in cyber D&D technology are also important, such as simulated intrusions and response.

Tracking honeytokens files is helpful in the monitoring stage, and can involve watermarking to alert defenders to an intruder. Correlating activity across disparate tripwires may indicate higher-level operational or strategic moves by an adversary.

Technical and operational metrics are needed to continuously improve cyber D&D operations in the *reinforce* stage. These measure the precision and believability of honeytokens in that they attract the intended target and are readily mistaken as real. Operational metrics measure how well a particular cyber D&D technique works with respect to a specific adversary and objective. The mark of an effective cyber D&D technique is its ability to entice the desired target and exploit its interests, preventing the adversary from compromising sensitive information.

Implementers can also measure the extent to which deception interferes with legitimate users, as well as the cost of deployment. At a macro level, organizations also need metrics for maintenance and development costs of cyber D&D operations. An efficient D&D technique is one that can be commoditized and produced in a modular and rapid manner.

THE MOST CRITICAL FACETS OF CYBER D&D ARE THE MATURITY LEVEL AND THE OVERALL MANAGEMENT MODEL.

RESEARCH AND TECHNOLOGY CHALLENGES

If enterprises move from a traditional cyber defense approach to one that incorporates cyber D&D, what are the grand challenges and hard technical problems that need to be solved?

domain-specific operators, but could be automatically created from a template. Related areas include tear-down and repurposing of honeytokens.

Historical analysis of intrusion attempts will inform the *plan* stage. What moves has the adversary made?

CALL FOR ACTION: STRATEGIC COOPERATION

Paradigms for cyber defense are evolving, from static and passive perimeter defenses to outward-looking active defenses that engage and study the adversary. This evolution opens the door at tactical, operational, and strategic levels for cyber D&D to enhance defenses in the national cyber strategy. Cyber D&D is an emerging interdisciplinary approach to cybersecurity, combining human research, experiments, and exercises into cyber D&D operations with the engineering and testing of appropriate cyber D&D TTTPs.

There is currently no national “center of gravity” for innovative research, standards development, shared repositories, or training curriculum creation for cyber D&D. Integration is lacking across the whole of government for policies and programs in cyber D&D, and for coordinating the development and use of cyber D&D defenses. Such a center of gravity would involve three action areas: standards, methodologies, and shared repositories; research and operational coordination; and active defense cyber D&D enterprise organization.

The first area focuses on best practices and standards for cyber D&D:

- › cataloging ongoing offensive and defensive cyber D&D techniques;
- › mapping ongoing threats to appropriate D&D techniques to support cyber intelligence on adversaries and intrusion campaign analysis;
- › conducting outcome analysis of operational cyber D&D techniques, impacts, and effectiveness; and
- › enabling the sharing of standards and methodologies

ABOUT THE AUTHORS

KRISTIN E. HECKMAN is a principal scientist and department head for artificial intelligence and cognitive science at the MITRE Corporation. Her research interests include denial and deception, cybersecurity, and neuropsychology. Heckman received a DSc in computer science from George Washington University. Contact her at kheckman@mitre.org.

FRANK J. STECH is a chief social scientist at the MITRE Corporation. His research interests include evidence-based intelligence analytics, communicated influence, and counterdeception. Stech received a PhD in psychology from the University of California, Santa Barbara. Contact him at stech@mitre.org.

BEN S. SCHMOKER is a security researcher at the MITRE Corporation. His research interests include malware analysis, intrusion detection, and exploit mitigation. Schmoker received a BS in computer science from the University of Nebraska Omaha. Contact him at bschmoker@mitre.org.

ROSHAN K. THOMAS is a principal security researcher at the MITRE Corporation. His research interests include dissemination and access controls, trust modeling, and modeling advanced persistent threat attacks. Thomas received a PhD in information technology from George Mason University. Contact him at rktomas@mitre.org.

through repositories of tools and practices to counter cyber threats with defensive D&D.

The second area focuses on facilitating four types of information exchange:

- › strategic, to formulate cyber D&D policy, programs, and sponsorship for participants and stakeholders;
- › research management, to formulate a strategic cyber D&D research roadmap with cyber researchers and operational community participation;

- › research, to share technical cyber D&D research; and
- › transformation, to formulate an operational roadmap that incorporates cyber D&D research results into cyber defense operations.

Government-sponsored research must lead the effort and commercial technology needs to make substantial investments.

Success in the third area requires an organization to serve as the trusted intermediary to broker cyber D&D operational sharing, collaboration, and networking to manage cyber D&D

defenses in the threat landscape. This organization would also organize collaboration of cyber D&D information exchange at highly detailed technical levels via technical exchange meetings, shared repositories, and standards. The organization would support the identification of near-term and long-term research needs and threat-based defense gaps, and help the whole of government to meet those needs. Finally, the organization would foster the development of cyber D&D training and educational curricula.

Cyber D&D should be part of the national cyber strategy. To achieve this goal, the national center of gravity program must facilitate a strategic “working group” to begin developing national cyber D&D plans, formulate US government policies, create programs,

and establish goals and objectives within the strategy. ■

REFERENCES

1. E. Waltz and M. Bennett, *Counter-deception Principles and Applications for National Security*, Artech House, 2007.
2. E.M. Hutchins, M.J. Cloppert, and R.M. Amin, “Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Cyber Kill Chains,” presentation at the 6th Ann. Int’l Conf. Information Warfare and Security, 2011; www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf.
3. B. Whaley, “Toward a General Theory of Deception,” J. Gooch and A. Perlmutter, eds., *Military Deception and Strategic Surprise*, Routledge, 2007, pp. 188–190.
4. The MITRE Corp., *Threat-Based Defense: A New Cyber Defense Playbook*, 2012; www.mitre.org/sites/default/files/pdf/cyber_defense_playbook.pdf.
5. Mandiant, *2014 Threat Report*; https://dl.mandiant.com/EE/library/WP_M-Trends2014_140409.pdf.
6. Mandiant, *APT1: Exposing One of China’s Cyber Espionage Units*; http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf.



See www.computer.org/computer-multimedia for multimedia content related to this article.



Call for Articles

IEEE Software seeks practical, readable articles that will appeal to experts and nonexperts alike.

The magazine aims to deliver reliable information to software developers and managers to help them stay on top of rapid technology change.

**IEEE
Software**

Author guidelines:

www.computer.org/software/author.htm
Further details: software@computer.org
www.computer.org/software